AF/S

TRANSMITTAL OF APPEAL BRIEF (Large Entity)

Docket No. POU920000176US1

In Re Application Of: **J. P. Kubala et al**

Application No.	Filing Date	Examiner	Customer No.	Group Art Unit	Confirmation No.
09/677,338	10/2/2000	Nilesh R. Shah	33558	2195	3544

Invention: **Method and Apparatus For Enforcing Capacity Limitations In A Logically Partitioned System**

COMMISSIONER FOR PATENTS:

Transmitted herewith in triplicate is the Appeal Brief in this application, with respect to the Notice of Appeal filed on **February 23, 2006**

The fee for filing this Appeal Brief is: **$500.00**

☐ A check in the amount of the fee is enclosed.

☒ The Director has already been authorized to charge fees in this application to a Deposit Account.

☒ The Director is hereby authorized to charge any fees which may be required, or credit any overpayment to Deposit Account No. **09-463**

☐ Payment by credit card. Form PTO-2038 is attached.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

Signature

Dated: **March 23, 2006**

William A. Kinnaman, Jr. - Attorney
Registration No, 27,650
2455 South Road - P386
Poughkeepsie, NY 12601

Telephone No. (845) 433-1175
Fax No. (845) 432-9601

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to "Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450" [37 CFR 1.8(a)] on

3/23/2006
(Date)

Sandra L. Kilmer
Signature of Person Mailing Correspondence

Sandra L. Kilmer
Typed or Printed Name of Person Mailing Correspondence

CC:

Patent

IN THE U.S. PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APEALS AND INTERFERENCES

Applicant: JEFFREY P. KUBALA et al.	:	Group Art Unit: 2195
Serial No.: 09/677,338	:	Examiner: Nilesh R. Shah
Filed: October 2, 2000	:	March 23, 2006
Confirmation No.: 3544	:	William A. Kinnaman, Jr.
Title: METHOD AND APPARATUS FOR	:	International Business Machines Corporation
ENFORCING CAPACITY LIMITATIONS	:	2455 South Road, Mail Station P386
IN A LOGICALLY PARTITIONED	:	Poughkeepsie, NY 12601
SYSTEM		

APPLICANTS' APPEAL BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Applicants hereby submit their appeal brief in the above-identified application.

CERTIFICATE OF MAILING UNDER 37 CFR 1.8(a)

I hereby certify that this correspondence is being deposited with the United States Postal Service as first-class mail in an envelope addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on March 23, 2006.

Sandra L. Kilmer — Date: 3-23-2006

REAL PARTY IN INTEREST

The real party in interest is International Business Machines Corporation, the assignee of record.

RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

STATUS OF CLAIMS

Claims 4-27 stand rejected and are on appeal. Claims 1-3 have been cancelled. No claims have been allowed or withdrawn from consideration.

STATUS OF AMENDMENTS

In an amendment after appeal filed February 24, 2006, claims 11, 20 and 26, directed to program storage devices, were rewritten as independent claims rather than as claims dependent on corresponding method claims. To date, nothing has been received regarding the status of this amendment.

Applicants are reproducing the claims in the Claims Appendix in the form in which they existed before this amendment. Whatever the disposition of the amendment, it should not affect the issues on appeal, as the amendment goes merely to the form of claims 11, 20 and 26 and not their substantive scope.

SUMMARY OF CLAIMED SUBJECT MATTER

Claims 4-15

Claim 4, which is representative of this first group of claims on appeal, is directed to a method of enforcing capacity limitations in an information handling system (Fig. 1: 100) in which a

physical machine (102) contains one or more logical partitions (108), each of which is allocated a defined portion of machine resources (104) and has one or more software applications (114) executing therein.

In accordance with this aspect of the invention, a maximum allowed consumption (Fig. 2A: C_{avg}) of the resources by one of the logical partitions is specified, and the actual consumption of the resources by the logical partition is measured. The actual consumption is compared with the maximum allowed consumption to determine whether it exceeds the maximum allowed consumption (Fig. 5C: step 518). If so, the actual consumption of the resources is reduced to the maximum allowed consumption by reducing the defined portion of machine resources allocated to the logical partition while allowing all of the software applications executing in the logical partition to continue executing (Fig. 5C: step 522).

Claim 11 is similar to claim 4, but is directed to a program storage device. Claim 12 is similar to claim 4, but is directed to apparatus.

Claims 16-24

Claim 16, which is representative of this group of claims on appeal, is similarly directed to a method of enforcing capacity limitations in an information handling system (100) in which a physical machine (102) contains one or more logical partitions (108), each of which is allocated a defined portion of machine resources (104).

In accordance with this aspect of the invention, a maximum capped consumption (Fig. 2A: C_{capped}) of the resources by one of the logical partitions is specified, and an actual average consumption of the resources by the logical partition is measured as before. The actual average consumption is compared with a maximum average consumption (Fig. 2A: C_{avg}) greater than the maximum capped consumption to determine whether the former exceeds the latter (Fig. 5C: step 518). If so, the actual average consumption of the resources is reduced to the maximum average consumption by alternatingly operating the logical partition in a capped mode (Fig. 2B: step 252) in which the logical partition is limited to the maximum capped consumption and in an uncapped

mode (Fig. 2B, step 254) in which the logical partition is not limited to the maximum capped consumption (Fig. 5C: step 522).

Claim 20 is similar to claim 16 but is directed to a program storage device. Claim 21 is similar to claim 16 but is directed to apparatus.

Claims 25-27

Claim 25, which is representative of this group of claims on appeal, is likewise directed to a method of enforcing capacity limitations in an information handling system (100) in which a physical machine (102) contains one or more logical partitions (108), each of which is allocated a defined portion of machine resources (104).

In accordance with this aspect of the invention, a maximum capped consumption (page 16, lines 26-27: C_{capped}) of the resources by one of a group of logical partitions is specified as a proportion (page 16, lines 26-27: [Partition_weight / (Σ All_active_partition_weights + Phantom_weight)]) of the resources (page 16, lines 26-27: Capacity) available to the group of logical partitions. Each logical partition in the group is assigned a weight (page 16, lines 26-27: Partition_weight), with the proportion being defined as the ratio of the weight of the logical partition to the sum of the weights of the logical partitions in the group. The logical partition is also assigned a phantom weight (page 16, lines 26-27: Phantom_weight) that is added to the sum (page 16, lines 26-27: Σ All_active_partition_weights) of the weights of the logical partitions in the group but not to the weight of the logical partition in determining that ratio.

An actual average consumption of the resources by the logical partition is measured and is compared with a maximum average consumption (Fig. 2A: C_{avg}) to determine whether the actual consumption exceeds the maximum average consumption. If it does, the actual average consumption of the resources is reduced to the maximum average consumption by operating the logical partition at least part of the time in a capped mode (Fig. 2B: step 252) in which the logical partition is limited to the maximum capped consumption (Fig. 5C: step 522).

Claim 26 is similar to claim 25, but is directed to a program storage device. Claim 27 is similar to claim 25, but is directed to apparatus.

GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

I. Claims 4-27 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Harris et al. U.S. Patent 6,438,704 ("Harris") in view of Beelitz U.S. Patent 6,032,239 ("Beelitz").[1]

ARGUMENT

Applicants will separately argue the groups of claims identified below.

Claims 4-15

Claim 4, which is representative of this group of claims, reads as follows:

> 4. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources and has one or more software applications executing therein, a method of enforcing capacity limitations comprising the steps of:
>
> specifying a maximum allowed consumption of said resources by one of said one or more logical partitions;
>
> measuring an actual consumption of said resources by said logical partition;
>
> comparing said actual consumption with said maximum allowed consumption to determine whether said actual consumption exceeds said maximum allowed consumption; and
>
> if said actual consumption exceeds said maximum allowed consumption, reducing said actual consumption of said resources to said maximum allowed consumption by reducing the defined portion of machine resources allocated to said logical partition while allowing all of the one or more software applications executing in said logical partition to continue executing.

Claim 4 thus recites (as do the remainder of claims 4-15) that if the actual consumption of resources by a logical partition exceeds the maximum allowed consumption, the actual consumption is reduced to the maximum allowed consumption by reducing the defined portion

[1] Office action of October 31, 2005 (hereinafter "Final Action"), page 2, ¶ 3.

of machine resources allocated to the logical partition while allowing the software applications executing in the logical partition to continue executing.

Claims 4-15, stand rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent 6,438,704 to Harris et al. ("Harris") in view of U.S. Patent 6,032,239 to Beelitz ("Beelitz") (Final Action, page 2, ¶ 3).

Harris relates to a system and method for scheduling use of system resources among a plurality of users. A share of processor time is specified for each user, and a particular user's CPU usage is limited to an absolute value in a "dispatch driven" multiprocessing system that is described in further detail in the patent. Among the mechanisms used by Harris is a "limit list" 27 (Fig. 1) of users whose CPU resource must be limited so that their consumption does not exceed a limit value; such users are said to have a "limit status" (col. 4, lines 9-12). Users on the limit list may be removed from the list and transferred to a dispatch list 37 (Fig. 1) to restore them to "dispatch status" when resources become available, as in low-load situations (col. 4, lines 12-16).

Beelitz relates to a system and method for updating partition mappings to logical drives in a computer memory device.

The Examiner contends that while Harris does not teach the use of different partitions, Beelitz teaches the use of different partitions with different resources "for upgrading, adding, deleting or changing the partition or resources on an exiting hard drive" (Final Action, page 3, ¶ 5). The Examiner concludes that it would have been obvious to combine the teachings of the two "because Beelitz's system of having changeable partitions would improve the efficiency of Harris' system by eliminating the particular partition that has hit its permitted consumption before it bogs down the entire system" (Final Action, page 3, ¶ 6). Applicants respectfully disagree.

As a threshold matter, while Beelitz relates to partitions, the partitions in question are those of a hard drive (Fig. 1: 16), where applications may reside before being loaded into main memory and executed. This has nothing to do with the partitioning of the physical machine (i.e., the CPU

and related resources) on which the applications are actually executed. Beelitz thus does not disclose a physical machine comprising one or more logical partitions, each of which is allocated a defined portion of machine resources "and has one or more software applications executing therein" as claimed by applicants.[2]

Of course, machines that are partitioned in the sense claimed by applicant exist in the prior art; applicants refer to several such machines in the background part of the specification, on page 2 at lines 8-14. However, even if one were to combine these teachings with those of Harris, one would still not have applicants' claimed invention.

From the Examiner's characterization of applicant's invention and Harris, it would seem that applicants' invention is nothing more than the system of Harris, but applied to a partitioned system. This is simply not so. In Harris, the resource consumption that is being managed is the consumption of machine resources (in particular, CPU time) by individual users. In applicants' claimed invention, it is the logical partition itself whose resource consumption is being measured and held to a maximum allowed consumption. Consumption of resources by individual users is irrelevant to this exercise unless, of course, that consumption pushes the total consumption for the partition over the limit.

Nothing in either Harris or the art relating to logical partitioning would suggest applying the monitoring techniques of Harris to a partition as whole rather than to an individual user running in that partition. This is particularly evident when one considers the motivation present in the two cases. In Harris, CPU usage by a particular user is limited in order to provide a "fair" amount of CPU time to other users competing for the same total CPU time (col. 1, lines 36-42). That is to say, Harris is seeking to manage the relative consumption of CPU resources by different users rather than the total CPU resource consumption. In applicants' system, on the other hand, the total resource consumption by a particular partition is limited in order to accommodate software

[2] The Examiner disputes applicants' assertion, contending that Beelitz "clearly teaches" a physical machine comprising one or more logical partitions, each of which is allocated a defined portion of machine resources "and has one or more software applications executing" (Final Action, page 6, ¶ 23). However, while Beelitz may have one or more applications "executing", those applications most first be loaded into main memory to be executed and are thus not executing in the partitioned resource (the hard drive) as claimed by applicants. For the same reason,

licensing schemes[3] that base the licensing fee on the total capacity available to the partition rather than the capacity actually used by the application (page 5, lines 4-28). Nothing in either of the references cited by the Examiner teaches managing the actual consumption of resources by a logical partition, as distinguished from individual users in a partition, in the manner claimed by applicants.

Accordingly, the Examiner's rejection of claims 4-15 as being unpatentable over these two references is untenable and should be reversed.

Claims 16-24

Claim 16, which is representative of this group of claims on appeal, reads as follows:

> 16. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources, a method of enforcing capacity limitations comprising the steps of:
> specifying a maximum capped consumption of said resources by one of said one or more logical partitions;
> measuring an actual average consumption of said resources by said logical partition;
> comparing said actual average consumption with a maximum average consumption greater than said maximum capped consumption to determine whether said actual average consumption exceeds said maximum average consumption; and
> if said actual average consumption exceeds said maximum average consumption, reducing said actual average consumption of said resources to said maximum average consumption by alternatingly operating said logical partition in a capped mode in which said logical partition is limited to said maximum capped consumption and in an uncapped mode in which said logical partition is not limited to said maximum capped consumption.

Claim 16 thus recites (as do the remainder of claims 16-24) that if the actual average consumption of resources by a logical partition exceeds a maximum average consumption, the actual average consumption is reduced to the maximum average consumption by alternatingly operating the logical partition in a capped mode in which the logical partition is limited to a

---

Beelitz's hard drive partitions cannot "run one or more pieces of software" as also asserted by the Examiner (Final Action, page 6, ¶ 23).

maximum capped consumption and in an uncapped mode in which the logical partition is not limited to the maximum capped consumption.

Since claims 16-24 incorporate the limitations of claims 4-15 relating to controlling resource consumption by a logical partition as a whole, they are believed to distinguish patentably over the references cited for reasons applicable to that other set of claims. They are additionally believed to distinguish patentably over the references cited by virtue of their recitations relating to alternatingly operating in a capped and in an uncapped mode.

Claims 16-24 likewise stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Harris in view of Beelitz (Final Action, page 2, ¶ 3).

In applying Harris, the Examiner is apparently equating being on the limit list 27 (Fig. 1) with applicants' capped mode and coming off the limit list with applicants' claimed uncapped mode of operation. This comparison is misplaced for several reasons.

First of all, as a general proposition, Harris is inapposite to applicants' claimed invention for the reasons given above, since it relates to controlling resource consumption by individual users, whereas applicants seek to control resource consumption by a logical partition, which may contain many such users. Nothing in either of the references cites suggests the desirability of controlling resource consumption by an entire logical partition.

Furthermore, Harris does not control the resource consumption of even individual users in the manner claimed by applicants for logical partitions. In applicants' claimed system, the logical partition is alternatingly operated in a capped mode in which it is limited to a maximum capped consumption (which is less than the maximum average consumption) and in an uncapped mode in which it is not limited to the maximum capped consumption, as shown in Fig. 2A. In either mode of operation, the partition continues to operate; it simply gets fewer (and capped) resources in capped mode than it does in uncapped mode. This contrasts with Harris' operation, in which

[3] Applicants' claims, to be sure, do not recite the use of such a software licensing scheme. Whether or not recited in applicants' claims, however, the actual motivation for applicants' claimed invention is relevant to the issue of obviousness.

"hard limit" users that have exceeded their limit are simply moved to the limit list and not run at all until their resource consumption comes within acceptable bounds (col. 2, lines 36-38). While this may be an acceptable mode of operation for individual users in the Harris system, it would not be an acceptable manner of controlling usage by an entire partition, since it would mean that no user would be able to run until the partition moved off the limit list.

Accordingly, the Examiner's rejection of claims 16-24 over the art cited is likewise untenable and should be reversed.

Claims 25-27

Claims 25, which is representative of this group of claims on appeal, reads as follows:

> 25. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources, a method of enforcing capacity limitations comprising the steps of:
>
> specifying a maximum capped consumption of said resources by one of said one or more logical partitions, said logical partition being one of a group of one or more logical partitions, said maximum capped consumption being specified as a proportion of the resources available to said group of logical partitions, each of the logical partitions in said group being assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group, said logical partition also being assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group but not to the weight of said logical partition in determining said ratio;
>
> measuring an actual average consumption of said resources by said logical partition;
>
> comparing said actual average consumption with a maximum average consumption to determine whether said actual consumption exceeds said maximum average consumption; and
>
> if said actual average consumption exceeds said maximum average consumption, reducing said actual average consumption of said resources to said maximum average consumption by operating said logical partition at least part of the time in a capped mode in which said logical partition is limited to said maximum capped consumption.

Claim 25 thus recites (as do claim 26-27) that the logical partition is one of a group of partitions, each of which is assigned a weight, that the maximum capped consumption is specified as a proportion of the resources available to the group of partitions, and that the proportion is defined

as the ratio of the weight of the logical partition to the sum of the weights of the partitions in the group. Additionally, the logical partition is assigned a phantom weight that is added to the sum of the weights of the logical partitions in the group but not to the weight of the partition itself in determining that ratio.

Since claims 25-27 incorporate the limitations of claims 4-15 relating to controlling resource consumption by a logical partition as a whole, as well the limitations of claims 16-24 relating to operating in a capped mode, they are believed to distinguish patentably over the references cited for reasons applicable to those other sets of claims. They are additionally believed to distinguish patentably over the references cited by virtue of their recitations relating to phantom weight.

Claims 25-27 likewise stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Harris in view of Beelitz (page 2, ¶ 3).

The Examiner takes official notice that uses of weights within partitions are "well known in the art" (page 4, ¶ 9). Such weights, however, are used to apportion CPU time or other system resources among various resource users on an unequally weighted basis. In a well-known manner, when determining the relative weight accorded to a given user, one divides the absolute weight assigned to that user by the sum of the absolute weights assigned to all users. In such a scenario, since each term in the denominator represents the weight of a particular user, the sum of the relative weights accorded to all users is by definition 100%. For example, if three users had absolute weights of 1, 2 and 5, their relative weights would be 12.5%, 25% and 67.5%, for a total of 100%.

This is not what happens, though, when assigning a phantom weight to a partition. The phantom weight does not represent the weight of the partition in question or any other partition, but is merely added to the denominator to reduce the relative weight of the partition. This concept of phantom weighting can even be applied when the machine in question has only a single logical partition; indeed, the concept is especially useful here, since otherwise there would be no ready mechanism for reducing the resource consumption by that partition to a specified maximum (page 14, lines 22-29).

Thus, while there may be considerable art on weighting generally, the Examiner can point to no art teaching the concept of a phantom weight, which is added to the sum of the partition weights in the denominator but does not represent the weight of any actual partition and thus is not added to the weight of the partition itself in the numerator to determine the ratio of the partition weight to the total weight. Accordingly, claims 8 and 25-27 distinguish patentably over the cited art by virtue of this claimed feature, as well as for the reasons urged above with respect to claims 4-15.

Claims 8 and 15

Claim 8, dependent (through claims 7 and 5) on claim 4 in the first group above, recites that a logical partition is assigned a phantom weight that is added to the sum of the weights of the logical partitions in a group in determining the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in the group. Claim 15 contains a similar recitation but is dependent (though claims 14 and 13) on corresponding apparatus claim 12.

Claims 8 and 15 are believed to distinguish patentably over the art cited for reasons applicable to the first group above (claims 4-15) as well as for reasons applicable to the third group above (claims 25-27).

Claims 19 and 24

Claim 19, dependent (through claims 18 and 17) on claim 16 in the second group above, recites that a logical partition is assigned a phantom weight that is added to the sum of the weights of the logical partitions in a group in determining the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in the group. Claim 24 contains a similar recitation but is dependent (though claims 23 and 22) on corresponding apparatus claim 21.

Claims 19 and 24 are believed to distinguish patentably over the art cited for reasons applicable to the second group above (claims 16-24) as well as for reasons applicable to the third group above (claims 25-27).

Conclusion

For the foregoing reasons, claims 4-27 clearly distinguish patentably over the art cited by the Examiner. Accordingly, the Examiner's rejection of these claims as being unpatentable over this art is untenable and should be reversed.

Respectfully submitted,
JEFFREY P. KUBALA et al.

By 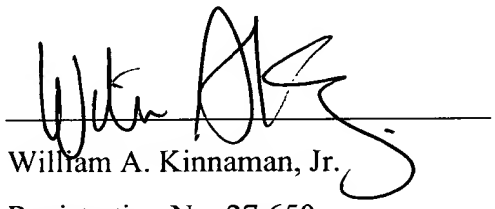

William A. Kinnaman, Jr.
Registration No. 27,650
Phone: (845) 433-1175
Fax: (845) 432-9601

WAK/wak

CLAIMS APPENDIX

4. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources and has one or more software applications executing therein, a method of enforcing capacity limitations comprising the steps of:

specifying a maximum allowed consumption of said resources by one of said one or more logical partitions;

measuring an actual consumption of said resources by said logical partition;

comparing said actual consumption with said maximum allowed consumption to determine whether said actual consumption exceeds said maximum allowed consumption; and

if said actual consumption exceeds said maximum allowed consumption, reducing said actual consumption of said resources to said maximum allowed consumption by reducing the defined portion of machine resources allocated to said logical partition while allowing all of the one or more software applications executing in said logical partition to continue executing.

5. The method of claim 4 in which said logical partition is one of a group of logical partitions, said maximum allowed consumption being specified as a proportion of the resources available to said group of logical partitions.

6. The method of claim 4 in which said group of logical partitions comprises all of the logical partitions on said physical machine.

7. The method of claim 5 in which each of the logical partitions in said group is assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group.

8. The method of claim 7 in which said logical partition is also assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group but not to the weight of said logical partition in determining said ratio.

9. The method of claim 4 in which said specified system resources are processor resources.

10. The method of claim 4 in which said actual consumption is determined as an rolling average of said consumption over a predetermined time interval.

11. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of claim 1.

12. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources and has one or more software applications executing therein, apparatus for enforcing capacity limitations comprising:

means for specifying a maximum allowed consumption of said resources by one of said one or more logical partitions;

means for measuring an actual consumption of said resources by said logical partition;

means for comparing said actual consumption with said maximum allowed consumption to determine whether said actual consumption exceeds said maximum allowed consumption; and

means for reducing said actual consumption of said resources to said maximum allowed consumption if said actual consumption exceeds said maximum allowed consumption by reducing the defined portion of machine resources allocated to said logical partition while allowing all of the one or more software applications executing in said logical partition to continue executing.

13. The apparatus of claim 12 in which said logical partition is one of a group of logical partitions, said maximum allowed consumption being specified as a proportion of the resources available to said group of logical partitions.

14. The apparatus of claim 13 in which each of the logical partitions in said group is assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group.

15. The apparatus of claim 14 in which said logical partition is also assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group in determining said ratio.

16. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources, a method of enforcing capacity limitations comprising the steps of:

specifying a maximum capped consumption of said resources by one of said one or more logical partitions;

measuring an actual average consumption of said resources by said logical partition;

comparing said actual average consumption with a maximum average consumption greater than said maximum capped consumption to determine whether said actual average consumption exceeds said maximum average consumption; and

if said actual average consumption exceeds said maximum average consumption, reducing said actual average consumption of said resources to said maximum average consumption by alternatingly operating said logical partition in a capped mode in which said logical partition is limited to said maximum capped consumption and in an uncapped mode in which said logical partition is not limited to said maximum capped consumption.

17. The method of claim 16 in which said logical partition is one of a group of one or more logical partitions, said maximum capped consumption being specified as a proportion of the resources available to said group of logical partitions.

18. The method of claim 17 in which each of the logical partitions in said group is assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group.

19. The method of claim 18 in which said logical partition is also assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group in determining said ratio.

20. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of claim 16.

21. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources, apparatus for enforcing capacity limitations comprising:

means for specifying a maximum capped consumption of said resources by one of said one or more logical partitions;

means for measuring an actual average consumption of said resources by said logical partition;

means for comparing said actual average consumption with a maximum average consumption greater than said maximum capped consumption to determine whether said actual average consumption exceeds said maximum average consumption; and

means for reducing said actual average consumption of said resources to said maximum average consumption if said actual average consumption exceeds said maximum average consumption by alternatingly operating said logical partition in a capped mode in which said logical partition is limited to said maximum capped consumption and in an uncapped mode in which said logical partition is not limited to said maximum capped consumption.

22. The apparatus of claim 21 in which said logical partition is one of a group of one or more logical partitions, said maximum capped consumption being specified as a proportion of the resources available to said group of logical partitions.

23. The apparatus of claim 22 in which each of the logical partitions in said group is assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group.

24. The apparatus of claim 23 in which said logical partition is also assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group in determining said ratio.

25. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources, a method of enforcing capacity limitations comprising the steps of:

specifying a maximum capped consumption of said resources by one of said one or more logical partitions, said logical partition being one of a group of one or more logical partitions, said maximum capped consumption being specified as a proportion of the resources available to said group of logical partitions, each of the logical partitions in said group being assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group, said logical partition also being assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group but not to the weight of said logical partition in determining said ratio;

measuring an actual average consumption of said resources by said logical partition;

comparing said actual average consumption with a maximum average consumption to determine whether said actual consumption exceeds said maximum average consumption; and

if said actual average consumption exceeds said maximum average consumption, reducing said actual average consumption of said resources to said maximum average consumption by operating said logical partition at least part of the time in a capped mode in which said logical partition is limited to said maximum capped consumption.

26. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of claim 25.

27. In an information handling system in which a physical machine contains one or more logical partitions, each of which is allocated a defined portion of machine resources, apparatus for enforcing capacity limitations comprising:

means for specifying a maximum capped consumption of said resources by one of said one or more logical partitions, said logical partition being one of a group of one or more logical partitions, said maximum capped consumption being specified as a proportion of the resources available to said group of logical partitions, each of the logical partitions in said group being assigned a weight, said proportion being defined as the ratio of the weight of said logical partition to the sum of the weights of the logical partitions in said group, said logical partition

also being assigned a phantom weight that is added to said sum of the weights of the logical partitions in said group but not to the weight of said logical partition in determining said ratio;

means for measuring an actual average consumption of said resources by said logical partition;

means comparing said actual average consumption with a maximum average consumption to determine whether said actual consumption exceeds said maximum average consumption; and

means for reducing said actual average consumption of said resources to said maximum average consumption if said actual average consumption exceeds said maximum average consumption by operating said logical partition at least part of the time in a capped mode in which said logical partition is limited to said maximum capped consumption.

EVIDENCE APPENDIX

(None)

RELATED PROCEEDINGS APPENDIX

(None)